

Fixed Point and Aperiodic Tilings

Bruno Durand¹, Andrei Romashchenko^{2,3}, Alexander Shen^{1,3}

¹ LIF, CNRS & Univ. de Provence, Marseille

² LIP, ENS de Lyon & CNRS

³ Institute for Information Transmission Problems, Moscow

Abstract. An aperiodic tile set was first constructed by R. Berger while proving the undecidability of the domino problem. It turned out that aperiodic tile sets appear in many topics ranging from logic (the Entscheidungsproblem) to physics (quasicrystals).

We present a new construction of an aperiodic tile set that is based on Kleene’s fixed-point construction instead of geometric arguments. This construction is similar to J. von Neumann self-reproducing automata; similar ideas were also used by P. Gács in the context of error-correcting computations.

The flexibility of this construction allows us to construct a “robust” aperiodic tile set that does not have periodic (or close to periodic) tilings even if we allow some (sparse enough) tiling errors. This property was not known for any of the existing aperiodic tile sets.

1 Introduction

In this paper⁴, *tiles* are unit squares with colored sides. Tiles are considered as prototypes: we may place translated copies of the same tile into different cells of

⁴ The first version of this preprint was published in `arxiv` and `hal` on 18 Feb 2008. Later this paper was published in proceedings of the DLT conference:

[DLT08] B.Durand, A.Romashchenko, A.Shen. Fixed Point and Aperiodic Tilings. *Proc. 12th international conference on Developments in Language Theory*. Kyoto, Japan, September 2008, pp. 537–548.

A short journal version of this work was presented in:

[EATCS] B.Durand, A.Romashchenko, A.Shen. Fixed point theorem and aperiodic tilings. *Bulletin of the EATCS (The Logic in Computer Science Column by Yuri Gurevich)*. no 97 (2009) pp. 126–136.

Also this article became a part of a long paper on a fixed-point technique in tilings:

[DRS09] B.Durand, A.Romashchenko, A.Shen. Fixed-point tile sets and their applications. 2009, `hal:00424024` and `arXiv:0910.2415` (50 pages).

Since the present paper is only a preliminary preprint, we encourage the reader to refer directly to [EATCS] or [DRS09] (this footnote is added on Jan 13, 2010).

a cell paper (rotations are not allowed). Tiles in the neighbor cells should match (common side should have the same color in both).

Formally speaking, we consider a finite set C of *colors*. A *tile* is a quadruple of colors (left, right, top and bottom ones), i.e., an element of C^4 . A *tile set* is a subset $\tau \subset C^4$. A *tiling* of the plane with tiles from τ (τ -*tiling*) is a mapping $U: \mathbb{Z}^2 \rightarrow \tau$ that respects the color matching condition. A tiling U is *periodic* if it has a *period*, i.e., a non-zero vector $T \in \mathbb{Z}^2$ such that $U(x + T) = U(x)$ for all $x \in \mathbb{Z}^2$. Otherwise the tiling is *aperiodic*. The following classical result was proved by Berger in a paper [2] where he used this construction as a main tool to prove *Berger's theorem*: the *domino problem* (to find out whether a given tile set has tilings or not) is undecidable.

Theorem 1. *There exists a tile set τ such that τ -tilings exist and all of them are aperiodic.* [2]

The first tile set of Berger was rather complicated. Later many other constructions were suggested. Some of them are simplified versions of the Berger's construction ([17], see also the expositions in [1,5,13]). Some others are based on polygonal tilings (including famous Penrose and Ammann tilings, see [10]). An ingenious construction suggested in [11] is based on the multiplication in a kind of positional number system and gives a small aperiodic set of 14 tiles (in [3] an improved version with 13 tiles is presented). Another nice construction with a short and simple proof (based explicitly on ideas of self-similarity) was recently proposed by N. Ollinger [16].

In this paper we present yet another construction of aperiodic tile set. It does not provide a small tile set; however, we find it interesting because:

- The existence of an aperiodic tile set becomes a simple application of a classical construction used in Kleene's fixed point (recursion) theorem, in von Neumann's self-reproducing automata [15] and, more recently, in Gács' reliable cellular automata [7,8]; we do not use any geometric tricks. The construction of an aperiodic tile set is not only an interesting result but an important tool (recall that it was invented to prove that domino problem is undecidable); our construction makes this tool easier to use (see Theorem 3).
- The construction is rather general, so it is flexible enough to achieve some additional properties of the tile set. Our main result is Theorem 6: there exists a "robust" aperiodic tile set that does not have periodic (or close to periodic) tilings even if we allow some (sparse enough) tiling errors. It is not clear whether this can be achieved for previously known aperiodic tile sets; however, the mathematical model for a processes like quasicrystals' growth or DNA-computation should take errors into account. Note that our model (independent choice of place where errors are allowed) has no direct physical meaning; it is just a simple mathematical model that can be used as a playground to develop tools for estimating the consequences of tiling errors.

The paper is organized as follows. In Section 2 we define the notion of a self-similar tile set (a tile set that simulates itself). In Section 3 we explain how a tile set can be simulated by a computation implemented by another tile set.

Section 4 shows how to achieve a fixed point (a tile set that simulates itself). Then we provide several applications of this construction: we use it to implement substitution rules (Section 5) and to obtain tile sets that are aperiodic in a strong sense (Section 6) and robust to tiling errors (Sections 7 and 8). Section 9 provides probability estimates that show that tiling errors are correctable with probability 1 (with respect to Bernoulli distribution). Finally, we show some other applications of the fixed point construction that simplify the proof of the undecidability of the domino problem and related results.

2 Macro-tiles

Fix a tile set τ and an integer $N > 1$. A *macro-tile* is an $N \times N$ square tiled by matching τ -tiles. Every side of a macro-tile carries a sequence of N colors called a *macro-color*.

Let ρ be a set of τ -macro-tiles. We say that τ *simulates* ρ if (a) τ -tilings exist, and (b) for every τ -tiling there exists a unique grid of vertical and horizontal lines that cuts this tiling into $N \times N$ macro-tiles from ρ .

Example 1. Assume that we have only one (‘white’) color and τ consists of a single tile with 4 white sides. Fix some N . There exists a single macro-tile of size $N \times N$. Let ρ be a singleton that contains this macro-tile. Then every τ -tiling can be cut into macro-tiles from ρ . However, τ does not simulate ρ , since the placement of cutting lines is not unique.




Fig. 1.

Example 2. In this example a set ρ that consists of exactly one macro-tile (that has the same macro-colors on all four sides) is simulated by some tile set τ . The tile set τ consists of N^2 tiles indexed by pairs (i, j) of integers modulo N . A tile from τ has colors on its sides as shown on Fig. 1. The macro-tile in ρ has colors $(0, 0), \dots, (0, N - 1)$ and $(0, 0), \dots, (N - 1, 0)$ on its borders (Fig. 2).

If a tile set τ simulates some set ρ of τ -macro-tiles with zoom factor $N > 1$ and ρ is isomorphic to τ , the set τ is called *self-similar*. Here an *isomorphism* between τ and ρ is a bijection that respects the relations “one tile can be placed on the right of another one” and “one tile can be placed on the top of another one”. (An isomorphism induces two bijections between horizontal/vertical colors of τ and horizontal/vertical macro-colors of ρ .)

The idea of self-similarity is used (more or less explicitly) in most constructions of aperiodic tile sets ([11,3] are exceptions); we find the following explicit formulation useful.




Fig. 2.

Theorem 2. A self-similar tile set τ has only aperiodic tilings.

Proof. Every τ -tiling U can be uniquely cut into $N \times N$ -macro-tiles from ρ . So every period T of U is a multiple of N (since the T -shift of a cut is also a

cut). Then T/N is a period of ρ -tiling, which is isomorphic to a τ -tiling, so T/N is again a multiple of N . Iterating this argument, we conclude that T is divisible by N^k for every k , so $T = 0$. \square

So to prove the existence of aperiodic tile sets it is enough to construct a self-similar tile set, and we construct it using the fixed-point idea. To achieve this, we first explain how to simulate a given tile set by embedding computations.

3 Simulating a tile set

For brevity we say that a tile set τ simulates a tile set ρ when τ simulates some set of macro tiles $\tilde{\rho}$ isomorphic to ρ (e.g., a self-similar tile set simulates itself).

Let us start with some informal discussion. Assume that we have a tile set ρ whose colors are k -bit strings ($C = \mathbb{B}^k$) and the set of tiles $\rho \subset C^4$ is presented as a predicate $R(c_1, c_2, c_3, c_4)$. Assume that we have some Turing machine \mathcal{R} that computes R . Let us show how to simulate ρ using some other tile set τ .

This construction extends Example 2, but simulates a tile set ρ that contains not a single tile but many tiles. We keep the coordinate system modulo N embedded into tiles of τ ; these coordinates guarantee that all τ -tilings can be uniquely cut into blocks of size $N \times N$ and every tile “knows” its position in the block (as in Example 2). In addition to the coordinate system, now each tile in τ carries supplementary colors (from a finite set specified below) on its sides. On the border of a macro-tile (i.e., when one of the coordinates is zero) only two supplementary colors (say, 0 and 1) are allowed. So the macro-color encodes a string of N bits (where N is the size of macro-tiles). We assume that $N \geq k$ and let k bits in the middle of macro-tile sides represent colors from C . All other bits on the sides are zeros (this is a restriction on tiles: each tile knows its coordinates so it also knows whether non-zero supplementary colors are allowed).

Now we need additional restrictions on tiles in τ that guarantee that the macro-colors on sides of each macro-tile satisfy the relation R . To achieve this, we ensure that bits from the macro-tile sides are transferred to the central part of the tile where the checking computation of \mathcal{R} is simulated (Fig. 3).




Fig. 3.

For that we need to fix which tiles in a macro-tile form “wires” (this can be done in any reasonable way; let us assume that wires do not cross each other) and then require that each of these tiles carries equal bits on two sides; again it is easy since each tile knows its coordinates.

Then we check R by a local rule that guarantees that the central part of a macro-tile represents a time-space diagram of \mathcal{R} ’s computation (the tape is horizontal, time goes up).

This is done in a standard way. We require that computation terminates in an accepting state: if not, the tiling cannot be formed.

To make this construction work, the size of macro-tile (N) should be large enough: we need enough space for k bits to propagate and enough time and space (=height and width) for all accepting computations of \mathcal{R} to terminate.

In this construction the number of supplementary colors depends on the machine \mathcal{R} (the more states it has, the more colors are needed in the computation zone). To avoid this dependency, we replace \mathcal{R} by a fixed universal Turing machine \mathcal{U} that runs a *program* simulating \mathcal{R} . Let us agree that the tape has an additional read-only layer. Each cell carries a bit that is not changed during the computation; these bits are used as a program for the universal machine (Fig. 4). So in the computation zone the columns carry unchanged bits, and the tile set restrictions guarantee that these bits form the program for \mathcal{U} , and the central zone represents the protocol of an accepting computation for that program. In this way we get a tile set τ that simulates ρ with zoom factor N using $O(N^2)$ tiles. (Again we need N to be large enough.)




Fig. 4.

4 Simulating itself

We know how to simulate a given tile set ρ (represented as a program for the universal TM) by another tile set τ with a large enough zoom factor N . Now we want τ to be isomorphic to ρ (then Theorem 2 guarantees aperiodicity). For this we use a construction that follows Kleene's recursion (fixed-point) theorem⁵ [12].

Note that most rules of τ do not depend on the program for \mathcal{R} , dealing with information transfer along the wires, the vertical propagation of unchanged program bits, and the space-time diagram for the universal TM in the computation zone. Making these rules a part of ρ 's definition (we let $k = 2 \log N + O(1)$ and encode $O(N^2)$ colors by $2 \log N + O(1)$ bits), we get a program that checks that macro-tiles behave like τ -tiles in this respect.

The only remaining part of the rules for τ is the hardwired program. We need to ensure that macro-tiles carry the same program as τ -tiles do. For that our program (for the universal TM) needs to access the bits of its own text. (This self-referential action is in fact quite legal: the program is written on the tape, and the machine can read it.) The program checks that if a macro-tile belongs to the first line of the computation zone, this macro-tile carries the correct bit of the program.

How should we choose N (hardwired in the program)? We need it to be large enough so the computation described (which deals with $O(\log N)$ bits) can fit

⁵ A reminder: Kleene's theorem says that for every transformation π of programs one can find a program p such that p and $\pi(p)$ produce the same output. Proof sketch: since the statement is language-independent (use translations in both directions before and after π), we may assume that the programming language has a function `GetText()` that returns the text of the program and a function `Exec(string s)` that replaces the current process by execution of a program s . (Think about an interpreter: surely it has an access to the program text; it can also recursively call itself with another program.) Then the fixed point is `Exec(π (GetText()))`.

in the computation zone. The computation is rather simple (polynomial in the input size, i.e., $O(\log N)$), so for large N it easily fits in $\Omega(N)$ available time.

This finishes the construction of a self-similar aperiodic tile set.

5 Substitution system and tilings

The construction of self-similar tiling is rather flexible and can be easily augmented to get a self-similar tiling with additional properties. Our first illustration is the simulation of substitution rules.

Let A be some finite alphabet and $m > 1$ be an integer. A *substitution rule* is a mapping $s: A \rightarrow A^{m \times m}$. By A -configuration we mean an integer lattice filled with letters from A , i.e., a mapping $\mathbb{Z}^2 \rightarrow A$ considered modulo translations.

A substitution rule s applied to a configuration X produces another configuration $s(X)$ where each letter $a \in A$ is replaced by an $m \times m$ matrix $s(a)$.

A configuration X is *compatible* with substitution rule s if there exists an infinite sequence $\dots \xrightarrow{s} X_3 \xrightarrow{s} X_2 \xrightarrow{s} X_1 \xrightarrow{s} X$, where X_i are some configurations.

Example 3. Let $A = \{0, 1\}$, $s(0) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $s(1) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. It is easy to see that the only configuration compatible with s is the chess-board coloring.

Example 4. Let $A = \{0, 1\}$, $s(0) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $s(1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$. One can check that all configurations that are compatible with this substitution rule (called *Thue – Morse configurations* in the sequel) are aperiodic.

The following theorem goes back to [14]. It says that every substitution rule can be enforced by a tile set.

Theorem 3 (Mozes). *Let A be an alphabet and let s be a substitution rule over A . Then there exists a tile set τ and a mapping $e: \tau \rightarrow A$ such that*

- (a) *s -image of any τ -tiling is an A -configuration compatible with s ;*
- (b) *every A -configuration compatible with s can be obtained in this way.*

Proof. We modify the construction of the tile set τ (with zoom factor N) taking s into account. Let us first consider the very special case when

- the substitution rule maps each A -letter into an $N \times N$ -matrix (i.e., $m = N$).
- the substitution rule is easy to compute: given a letter $u \in A$ and (i, j) , we can compute the (i, j) -th letter of $s(u)$ in time $\text{poly}(\log |A|) \ll N$.

In this case we proceed as follows. In our basic construction every tile knows its coordinates in the macro-tile and some additional information needed to arrange ‘wires’ and simulate calculations of the universal TM. Now in addition to this basic structure each tile keeps two letters of A : the first is the label of a tile itself, and the second is the label of the $N \times N$ -tile it belongs to. This means that we keep additional $2 \log |A|$ bits in each tile, i.e., multiply the number of tiles by $|A|^2$. It remains to explain how the local rules work. We add two requirements:

- (a) the second letter is the same for neighbor tiles (unless they are separated by a border of some $N \times N$ macro-tile);
- (b) the first letter in a tile is determined by the second letter and the coordinates of the tile inside the macro-tile, according to the substitution rule.

Both requirements are easy to integrate in our construction. The requirement (a) is rather trivial; to achieve (b) we need to embed in a macro-tile a calculation of $s(\text{[label on this macro-tile]})$. It is possible when s is easy to compute.

The requirements (a) and (b) ensure that configuration is an s -image of some other configuration. Also (due the self-similarity) we have the same at the level of macro-tiles. But this is not all: we need to guarantee that the first letter on the level of macro-tiles is identical to the second letter on the level of tiles. This is also achievable: the first letter of a macro-tile is encoded by bits on its border, and we can require that these bits match the second letter of the tiles at that place (recall that second letter is the same across the macro-tile). It is easy to see that now τ has the required properties (each tiling projects into a configuration compatible with τ and vice versa).

However, this construction assumes that N (the zoom factor) is equal to the matrix size in the substitution rule, which is usually not the case (m is given, and N we have to choose, and it needs to be large enough). The solution is to let N be equal to m^k for some k , and use the substitution rule s^k , i.e., the k -th iteration of s (a configuration is compatible with s^k if and only if it is compatible with s). Now we do not need s to be easily computed: for large k the computation of s^k will fit into the space available (exponential in k). \square

6 Strong version of aperiodicity

Let $\alpha > 0$ be a real number. A configuration $U: \mathbb{Z}^2 \rightarrow A$ is α -aperiodic if for every nonzero vector $T \in \mathbb{Z}^2$ there exists N such that in every square whose side is at least N the fraction of points x such that $U(x) \neq U(x + T)$ exceeds α .

Remark 1. If U is α -aperiodic, then Besicovitch distance between U and any periodic pattern is at least $\alpha/2$. (The Besicovitch distance is defined as $\limsup_N d_N$ where d_N is the fraction of points where two patterns differ in the $N \times N$ centered square.)

Theorem 4. *There exists a tile set τ such that τ -tilings exist and every τ -tiling is α -aperiodic for every $\alpha < 1/3$.*

Proof. This tile set is obtained by applying Theorem 3 to Thue–Morse substitution rule T (Example 4). Note that any configuration $C = \{c_{ij}\}$ compatible with T is a xor-combination $c_{ij} = a_i \oplus b_j$ of two one-dimensional Thue–Morse sequences a and b , and for a and b a similar result (every shift changes between $1/3$ and $2/3$ of positions in a large block) is well known (see, e.g., [18]). \square

7 Filling holes

The second application of our flexible fixed-point construction is an aperiodic

tile set where isolated defects can be healed.

Let $c_1 < c_2$ be positive integers. We say that a tile set τ is (c_1, c_2) -robust if the following holds: For every n and for every τ -tiling U of the c_2n -neighborhood of a square $n \times n$ excluding the square itself there exists a tiling V of the entire c_2n -neighborhood of the square (including the square itself) that coincides with U outside of the c_1n -neighborhood of the square (see Fig. 5).




Fig. 5.

Theorem 5. *There exists a self-similar tile set that is (c_1, c_2) -robust for some c_1 and c_2 .*

Proof. For every tile set μ it is easy to construct a “robustified” version μ' of μ , i.e., a tile set μ' and a mapping $\delta: \mu' \rightarrow \mu$ such that: (a) δ -images of μ' -tilings are exactly μ -tilings; (b) μ' is “5-robust”: every μ' -tiling of a 5×5 square minus 3×3 hole can be uniquely extended to the tiling of the entire 5×5 square.




Fig. 6.

Indeed, it is enough to keep in one μ' -tile the information about, say, 5×5 square in μ -tiling and use the colors on the borders to ensure that this information is consistent in neighbor tiles.

This robustification can be easily combined with the fixed-point construction. In this way we can get a 5-robust self-similar tile set τ if the zoom factors N is large enough. Let us show that this set is also (c_1, c_2) -robust for some c_1 and c_2 (that depend on N , but N is fixed.)

Indeed, let us have a tiling of a large enough neighborhood around an $n \times n$ hole. Denote by k the minimal integer such that $N^k \geq n$ (so the k -level macro-tiles are greater than the hole under consideration). Note that the size of k -level macro-tiles is linear in n since $N^k \leq N \cdot n$.

In the tiling around the hole, an $N \times N$ block structure is correct except for the N -neighborhood of the central $n \times n$ hole. For similar reasons $N^2 \times N^2$ -structure is correct except for the $N + N^2$ -neighborhood, etc. So for the chosen k we get a k -level structure that is correct except for (at most) $9 = 3 \times 3$ squares of level k , and such a hole can be filled (due to 5-robustness) with $N^k \times N^k$ squares, and these squares can be then detalized back.

To implement this procedure (and fill the hole), we need a correct tiling only in the $O(N^k)$ -neighborhood of the hole (technically, we need to have a correct tiling in $(3N^k)$ -neighborhood of the hole; as $3N^k \leq 3Nn$, we let $c_2 = 3N$). The correction procedure involves changes in another $O(N^k)$ -neighborhood of the hole (technically, changes touch $(2N^k)$ -neighborhood of the hole; $2N^k \leq 2Nn$, so we let $c_1 = 2N$). \square

8 Tilings with errors

Now we combine our tools to prove that there exists a tile set τ that is aperiodic in rather strong sense: this set does not have periodic tilings or tilings that are close to periodic. Moreover, this remains true if we allow the tiling to have some “sparse enough” set of errors. Tiling with errors is no more a tiling (as

defined above): in some places the neighbor colors do not match. Technically it is more convenient to consider tilings with “holes” (where some cells are not tiled) instead of errors but this does not matter: we can convert a tiling error into a hole just by deleting one of two non-matching tiles.

Let τ be a tile set and let $H \subset \mathbb{Z}^2$ be some set (H for “holes”). We consider (τ, H) -tilings, i.e., mappings $U: \mathbb{Z}^2 \setminus H \rightarrow \tau$ such that every two neighbor tiles from $\mathbb{Z}^2 \setminus H$ match (i.e., have the same color on the common side).

We claim that there exists a tile set τ such that (1) τ -tilings of the entire plane exist and (2) for every “sparse enough” set H every (τ, H) -tiling is far from every periodic mapping $\mathbb{Z}^2 \rightarrow \tau$.

To make this claim true, we need a proper definition of a “sparse” set. The following trivial counterexample shows that a requirement of small density is not enough for such a definition: if H is a grid made of vertical and horizontal lines at large distance N , the density of H is small but for any τ there exist (τ, H) -tilings with periods that are multiples of N .

The definition of sparsity we use (see below) is rather technical; however, it guarantees that for small enough ε a random set where every point appears with probability ε independently of other points, is sparse with probability 1. More precisely, for every $\varepsilon \in (0, 1)$ consider a Bernoulli probability distribution B_ε on subsets of \mathbb{Z}^2 where each point is included in the random subset with probability ε and different points are independent.

Theorem 6. *There exists a tile set τ with the following properties: (1) τ -tilings of \mathbb{Z}^2 exist; (2) for all sufficiently small ε for almost every (with respect to B_ε) subset $H \subset \mathbb{Z}^2$ every (τ, H) -tiling is at least $1/10$ Besicovitch-apart from every periodic mapping $\mathbb{Z}^2 \rightarrow \tau$.*

Remark 2. Since the tiling contains holes, we need to specify how we treat the holes when defining Besicovitch distance. We do *not* count points in H as points where two mappings differ; this makes our statement stronger.

Remark 3. The constant $1/10$ is not optimal and can be improved by a more accurate estimate.

Proof. Consider a tile set τ such that (a) all τ -tilings are α -aperiodic for every $\alpha < 1/3$; (b) τ is (c_1, c_2) -robust for some c_1 and c_2 . Such a tile set can be easily constructed by combining the arguments used for Theorem 5 and Theorem 4.

Then we show (this is the most technical part postponed until Section 9) that for small ε a B_ε -random set H with probability 1 has the following “error-correction” property: every (τ, H) -tiling is Besicovitch-close to some τ -tiling of the entire plane. The latter one is α -aperiodic, therefore (if Besicovitch distance is small compared to α) the initial (τ, H) -tiling is far from any periodic mapping.

For simple tile sets that allow only periodic tilings this error-correction property can be derived from basic results in percolation theory (the complement of H has large connected component etc.) However, for aperiodic tile sets this argument does not work and we need more complicated notion of “sparse” set based on “islands of errors”. We employ the technique suggested in [7] (see also applications of “islands of errors” in [9], [6]).

9 Islands of errors

Let $E \subset \mathbb{Z}^2$ be a set of points; points in E are called *dirty*; other points are *clean*. Let $\beta \geq \alpha > 0$ be integers. A set $X \subset E$ is an (α, β) -island in E if:

- (1) the diameter of X does not exceed α ;
- (2) in the β -neighborhood of X there is no other points from E .

(Diameter of a set is a maximal distance between its elements; the distance d is defined as the maximum of distances along both coordinates; β -neighborhood of X is a set of all points y such that $d(y, x) \leq \beta$ for some $x \in X$.)

It is easy to see that two (different) islands are disjoint (and the distance between their points is greater than β).

Let $(\alpha_1, \beta_1), (\alpha_2, \beta_2), \dots$ be a sequence of pairs of integers and $\alpha_i \leq \beta_i$ for all i . Consider the iterative “cleaning” procedure. At the first step we find all (α_1, β_1) -islands (*rank 1 islands*) and remove all their elements from E (thus getting a smaller set E_1). Then we find all (α_2, β_2) -islands in E_1 (*rank 2 islands*); removing them, we get $E_2 \subset E_1$, etc. Cleaning process is *successful* if every dirty point is removed at some stage.

At the i th step we also keep track of the β_i -neighborhoods of islands deleted during this step. A point $x \in \mathbb{Z}^2$ is *affected* during a step i if x belongs to one of these neighborhoods.

The set E is called *sparse* (for given sequence α_i, β_i) if the cleaning process is successful, and, moreover, every point $x \in \mathbb{Z}^2$ is affected at finitely many steps only (i.e., x is far from islands of large ranks).

The values of α_i and β_i should be chosen in such a way that:

- (1) for sufficiently small $\varepsilon > 0$ a B_ε -random set is sparse with probability 1 (Lemma 1 below);
- (2) if a tile set τ is (c_1, c_2) -robust and H is sparse, then any (τ, H) -tiling is Besicovitch close to some τ -tiling of the entire plane (Lemmas 2 and 3).

Lemma 1. Assume that $8 \sum_{k < n} \beta_k < \alpha_n \leq \beta_n$ for every n and $\sum_i \frac{\log \beta_i}{2^i} < \infty$. Then for all sufficiently small $\varepsilon > 0$ a B_ε -random set is sparse with probability 1.




Fig. 7. Explanation tree; vertical lines connect different names for the same points.

x in E_{n-1} is an (α_n, β_n) -island in E_{n-1} and x would be removed.

Each of the points x_1 and x (that we denote also x_0 to make the notation uniform) belongs to E_{n-1} because it belongs to E_{n-2} together with some other

Proof of Lemma 1. Let us estimate the probability of the event “ x is not cleaned after n steps” for a given point x (this probability does not depend on x). If $x \in E_n$, then x belongs to E_{n-1} and is not cleaned during the n th step (when (α_n, β_n) -islands in E_{n-1} are removed). Then $x \in E_{n-1}$ and, moreover, there exists some other point $x_1 \in E_{n-1}$ such that $d(x, x_1)$ is greater than $\alpha_n/2$ but not greater than $\beta_n + \alpha_n/2 < 2\beta_n$. Indeed, if there were no such x_1 in E_{n-1} , then $\alpha_n/2$ -neighborhood of

point (at the distance greater than $\alpha_{n-1}/2$ but not exceeding $\beta_{n-1} + \alpha_{n-1}/2$). In this way we get a tree (Figure 7) that explains why x belongs to E_n .

The distance between x_0 and x_1 in this tree is at least $\alpha_n/2$ while the diameter of the subtrees starting at x_0 and x_1 does not exceed $\sum_{i < n} 2\beta_i$. Therefore, the Lemma's assumption guarantees that these subtrees cannot intersect and, moreover, that all the leaves of the tree are different. Note that all 2^n leaves of the tree belong to $E = E_0$. As every point appears in E independently from other points, such an “explanation tree” is valid with probability ε^{2^n} . It remains to estimate the number of possible explanation trees for a given point x .

To specify x_1 we need to specify horizontal and vertical distance between x_0 and x_1 . Both distances do not exceed $2\beta_n$, therefore we need about $2 \log(4\beta_n)$ bits to specify them (including the sign bits). Then we need to specify the distances between x_{00} and x_{01} as well as distances between x_{10} and x_{11} ; this requires at most $4 \log(4\beta_{n-1})$ bits. To specify the entire tree we therefore need

$$2 \log(4\beta_n) + 4 \log(4\beta_{n-1}) + 8 \log(4\beta_{n-2}) + \dots + 2^n \log(4\beta_1),$$

that is (reversing the sum and taking out the factor 2^n) equal to $2^n(\log(4\beta_1) + \log(4\beta_2)/2 + \dots)$. Since the series $\sum \log \beta_n / 2^n$ converges by assumption, the total number of explanation trees for a given point (and given n) does not exceed $2^{O(2^n)}$, so the probability for a given point x to be in E_n for a B_ε -random E does not exceed $\varepsilon^{2^n} 2^{O(2^n)}$, which tends to 0 (even super-exponentially fast) as $n \rightarrow \infty$.

We conclude that the event “ x is not cleaned” (for a given point x) has zero probability; the countable additivity guarantees that with probability 1 all points in \mathbb{Z}^2 are cleaned.

It remains to show that every point with probability 1 is affected by finitely many steps only. Indeed, if x is affected by step n , then some point in its β_n -neighborhood belongs to E_n , and the probability of this event is at most $O(\beta_n^2) \varepsilon^{2^n} 2^{O(2^n)} = 2^{2 \log \beta_n + O(2^n) - \log(1/\varepsilon) 2^n}$; the convergence conditions guarantees that $\log \beta_n = o(2^n)$, so the first term is negligible compared to others, the probability series converges and the Borel–Cantelli lemma gives the desired result. \square

The following (almost evident) Lemma describes the error correction process.

Lemma 2. Assume that a tile set τ is (c_1, c_2) -robust, $\beta_k > 4c_2\alpha_k$ for every k and a set $H \subset \mathbb{Z}^2$ is sparse (with respect to α_i, β_i). Then every (τ, H) -tiling can be transformed into a τ -tiling of the entire plane by changing it in the union of $2c_1\alpha_k$ -neighborhoods of rank k islands (for all islands of all ranks).

Proof of Lemma 2. Note that $\beta_k/2$ -neighborhoods of rank k islands are disjoint and large enough to perform the error correction of rank k islands, since $\beta_k > 4c_2\alpha_k$. \square

It remains to estimate the Besicovitch size of the part of the plane changed during error correction.

Lemma 3. The Besicovitch distance between the original and corrected tilings (in Lemma 2) does not exceed $O(\sum_k (\alpha_k/\beta_k)^2)$. (Note that the constant in O -notation depends on c_1 .)

Proof of Lemma 3. We need to estimate the fraction of changed points in large centered squares. By assumption, the center is affected only by a finite number of islands. For every larger rank k , the fraction of points affected at the stage k in *any* centered square does not exceed $O((\alpha_k/\beta_k)^2)$: if the square intersects with the changed part, it includes a significant portion of the unchanged part. For smaller ranks the same is true for *all large enough* squares that cover completely the island affecting the center point). \square

It remains to choose α_k and β_k . We have to satisfy all the inequalities in Lemmas 1–3 at the same time. To satisfy Lemma 2 and Lemma 3, we may let $\beta_k = c\alpha_k$ for large enough c . To satisfy Lemma 1, we may let $\alpha_{k+1} = 8(\beta_1 + \dots + \beta_k) + 1$. Then α_k and β_k grow faster than any geometric sequence (like factorial multiplied by a geometric sequence), but still $\log \beta_i$ is bounded by a polynomial in i and the series in Lemma 1 converges.

With these parameters (taking c large enough) we may guarantee that Besicovitch distance between the original (τ, H) -tiling and the corrected τ -tiling does not exceed, say $1/100$. Since the corrected tiling is $1/5$ -aperiodic and $1/10 + 2 \cdot (1/100) < 1/5$, we get the desired result (Theorem 6). \square

10 Other applications of fixed point self-similar tilings

The fixed point construction of aperiodic tile set is flexible enough and can be used in other contexts. For example, the “zoom factor” N can depend on the level k (number of grouping steps). For this each macro-tile should have k encoded at its sides; this labeling should be consistent when switching to the next level. For a tile of level k its coordinates inside a macro-tile are integers modulo N_{k+1} , so in total $\log k + O(\log N_{k+1})$ bits are required and N_k steps should be enough to perform addition modulo N_{k+1} . This means that N_k should not increase too fast or too slow (say, $N_k = \log k$ is too slow and $N_{k+1} = 2^{N_k}$ is too fast). Also we need to compute N_k when k is known, so we assume that this can be done in polynomial time in the length of k (i.e., $\log k$). These restrictions still allow many possibilities, say, $N_k = \sqrt{k}$, $N_k = k$, $N_k = 2^{(2^k)}$, $N_k = k!$ etc.

This “self-similar” structure with variable zoom factor can be useful in some cases. Though it is not a self-similar according to our definition, one can still easily prove that any tiling is aperiodic. Note that now the computation time for the TM simulated in the central part increases with level, and this can be used for a simple proof of undecidability of domino problem (in the standard proof [2,1] one needs to organize the “computation zone” with some simple geometric tricks). With our new construction it is enough (for a given TM M) to add in the program the parallel computation of M on the empty tape; if it terminates, this destroys the tiling. This construction can be used to replace the constant $1/10$ in Theorem 6 by any number less than 1 ; to provide a new proof for the results of [4] (a tileset whose tilings have maximal Kolmogorov complexity) and extend them to tilings with sparse errors; it can be also used in some other applications of tilings. Here is another application of this construction. We say that a tile set τ is *m-periodic* if τ -tilings exist and for each of them the set of

periods is the set of *all* multiples of m (this is equivalent to the fact that both vectors $(0, m)$ and $(m, 0)$ are periods). Let E [resp. O] be all m -periodic tile sets for all even m [resp. odd m].

Theorem 7. *The sets E and O are inseparable enumerable sets.*

Acknowledgments. The authors thank the participants of the Kolmogorov seminar in Moscow (working on the RFBR projects 05-01-02803 and 06-01-00122-a) for many fruitful discussions.

References

1. C. Allauzen, B. Durand. Appendix A: Tiling Problems. In: E. Börger, E. Grädel, Y. Gurevich, *The Classical Decision Problems*, Springer-Verlag, 1996.
2. R. Berger, The Undecidability of the Domino Problem. *Mem. Amer. Math. Soc.*, **66**, 1966.
3. K. Culik, An Aperiodic Set of 13 Wang Tiles, *Discrete Math.*, **160**, 245–251, 1996.
4. B. Durand, L. Levin, A. Shen, Complex Tilings. *J. Symbolic Logic*, **73** (2), 593–613, 2008 (See also Proc. *33rd Ann. ACM Symp. Theory Computing*, 732–739, 2001, and www.arxiv.org/cs.CC/0107008 for an earlier version.)
5. B. Durand, L. Levin, A. Shen, Local Rules and Global Order, or Aperiodic Tilings, *Math. Intelligencer*, **27**(1), 64–68, 2004.
6. B. Durand, A. Romashchenko, On Stability of Computations by Cellular Automata, In Proc. *European Conf. Compl. Syst.*, Paris, 2005.
7. P. Gács, Reliable Cellular Automata with Self-Organization, In Proc. *38th Ann. Symp. Found. Comput. Sci.*, 90–97, 1997.
8. P. Gács, Reliable Cellular Automata with Self-Organization, *J. Stat. Phys.*, **103** (1/2), 45–267, 2001.
9. L. Gray, A Reader’s Guide to Gács’ Positive Rates Paper. *J. Stat. Phys.*, **103** (1/2), 1–44, 2001.
10. B. Grünbaum, G.C. Shephard, *Tilings and Patterns*, W.H. Freeman and Company, New York, 1987.
11. J. Kari, A Small Aperiodic Set of Wang tiles, *Discrete Math.*, **160**, 259–264, 1996.
12. H. Rogers, *The Theory of Recursive Functions and Effective Computability*, Cambridge, MIT Press, 1987.
13. L. Levin, Aperiodic Tilings: Breaking Translational Symmetry, *Computer J.*, **48**(6), 642–645, 2005. On-line: <http://www.arxiv.org/cs.DM/0409024>
14. S. Mozes, Tilings, Substitution Systems and Dynamical Systems Generated by Them, *J. Analyse Math.*, **53**, 139–186, 1989.
15. J. von Neumann, *Theory of Self-reproducing Automata*, Edited by A. Burks, University of Illinois Press, 1966.
16. N. Ollinger, *Two-by-two Substitution Systems and the Undecidability of the Domino Problem*, In Proc. *Computability in Europe*, LNCS **5028**, 476–485, 2008.
17. R. Robinson, Undecidability and Nonperiodicity for Tilings of the Plane. *Inventiones Mathematicae*, **12**, 177–209, 1971.
18. M. Zaks, A.S. Pikovsky, J. Kurths, On the Correlation Dimension of the Spectral Measure for the Thue–Morse Sequence, *J. Stat. Phys.*, **88**(5/6), 1387–1392, 1997.